

Oral Cancer Detection and Medical Diagnosis System Using Deep Learning Techniques

Narawade Vaishnavi Nitin, Narawade Bhushan Navnath, Prof. Sanchita Panmand

Master of Computer Applications.

Samarth College of Engineering & Management, Belhe

vaishnavinarawade76@gmail.com, bhushannarawade22@gmail.com, sanchitapanmand19@gmail.com

Abstract: *Oral cancer is a serious health condition that affects different parts of the oral cavity, including the tongue, gums, cheeks, and lips. Early identification of suspicious lesions is important for timely medical evaluation and treatment. This research paper presents an Oral Cancer Detection and Medical Diagnosis System Using Deep Learning Techniques to support preliminary oral cancer screening through automated image classification. The proposed system uses a Convolutional Neural Network (CNN) based on the MobileNetV2 architecture to classify oral images into benign and malignant categories. The model processes uploaded images after resizing them to (224×224) pixels and normalizing their pixel values. Python, TensorFlow/Keras, and Flask are used to integrate the deep learning model with a web-based application.*

The system provides user registration, login, oral image uploading, and prediction result display functionalities, while SQLite manages user account information. The proposed approach demonstrates the practical integration of artificial intelligence, image processing, and web technologies for medical image analysis. Its effectiveness depends on dataset quality, model training, and evaluation using independent test images. Performance measures such as accuracy, precision, recall, specificity, and F1-score should be calculated through actual experiments. The system is intended for preliminary screening support and academic research, not as a replacement for professional medical examination or confirmatory diagnosis. Future improvements may include dataset expansion, model performance evaluation, explainable artificial intelligence, and clinical validation.

Keywords: Oral Cancer Detection, Deep Learning, CNN, MobileNetV2, Medical Image Classification, Artificial Intelligence, Image Processing, Flask, TensorFlow, Oral Cancer Screening.

I. INTRODUCTION

Oral cancer is a major health concern that affects the tissues of the oral cavity, including the tongue, lips, gums, inner cheeks, and other areas of the mouth. Early identification of suspicious lesions is important because timely medical evaluation can improve the chances of effective treatment. Traditional oral cancer screening methods mainly involve visual examination by dentists and healthcare professionals, followed by additional investigations when necessary. However, identifying abnormal tissue changes can be challenging because lesions may differ in colour, shape, size, and texture. Variations in clinical experience and image quality can also affect the interpretation of oral lesions. Therefore, the application of advanced computational techniques to oral image analysis has become an important area of research. Artificial Intelligence (AI) and deep learning have created new opportunities for automated medical image analysis. Deep learning models can learn relevant visual features from images and use these features to perform classification tasks. Convolutional Neural Networks (CNNs) are widely used in image-processing applications because they can identify patterns such as edges, textures, shapes, and complex structures. In oral cancer screening research, these techniques can be investigated to distinguish between images belonging to benign and malignant categories. MobileNetV2 is a lightweight CNN architecture designed to provide computationally efficient image processing. Its



use in an oral image classification system offers an opportunity to explore deep learning techniques in a practical software application.

The proposed Oral Cancer Detection and Medical Diagnosis System Using Deep Learning Techniques integrates a trained MobileNetV2-based model with a web application developed using Python and Flask. The system accepts oral images, resizes them to (224×224) pixels, normalizes their pixel values, and passes the processed images to the trained model for classification. TensorFlow and Keras support model loading and prediction, while HTML, CSS, and JavaScript provide the user interface. SQLite is used to store user account information, and the application provides registration, login, image submission, and prediction result display functionalities. This integration demonstrates how deep learning and web technologies can be combined to develop an accessible image-classification application.

II. PROBLEM STATEMENT

Oral cancer is a serious disease, and its early detection is essential for timely treatment. Traditional detection methods depend on clinical examination and medical expertise, which may require time and access to specialized healthcare services. Therefore, this project, Oral Cancer Detection and Medical Diagnosis System Using Deep Learning Techniques, aims to develop a web-based system using CNN and MobileNetV2 to analyze uploaded oral images and classify them as benign or malignant. The system uses Python, TensorFlow/Keras, and Flask to provide predicted results and confidence scores, supporting preliminary screening. However, the system is intended only as an assistive tool and cannot replace professional medical evaluation or confirm a cancer diagnosis.

III. OBJECTIVES

- To develop a deep learning-based system for preliminary detection of oral cancer using oral images.
- To implement a CNN model based on MobileNetV2 for classifying oral images as benign or malignant.
- To provide a user-friendly web interface for uploading oral images and displaying prediction results.
- To display the predicted classification and confidence score for the uploaded image.
- To develop the system using Python, Flask, TensorFlow/Keras, and a suitable database for user management.

IV. LITERATURE SURVEY

1. Madhusmita Das, Rasmita Dash and Sambit Kumar Mishra (2023)

Das et al. (2023) proposed an automatic oral squamous cell carcinoma detection system using a Convolutional Neural Network (CNN). The researchers used histopathological images of oral mucosa to classify oral cancer and compared their approach with established deep learning architectures, including VGG16, VGG19, ResNet50, MobileNet and InceptionNet. Their research demonstrated the potential of CNN-based techniques for automated oral cancer image analysis and preliminary identification.

2. Santisudha Panigrahi, Bhabani Sankar Nanda and Ruchi Bhuyan (2023)

Panigrahi et al. (2023) studied the classification of oral squamous cell carcinoma using deep transfer learning techniques. Their research investigated pre-trained deep learning models for analyzing histopathological images and identifying cancer-related patterns. Transfer learning can reduce the need to train a deep neural network entirely from scratch and can help improve feature extraction. This study is relevant to the proposed project because it supports the use of deep learning models for oral cancer image classification.

3. Hitoshi Nagatsuka et al. (2023)

Nagatsuka et al. (2023) investigated the effectiveness of deep learning classifiers in the histopathological diagnosis of oral squamous cell carcinoma. The study examined CNN architectures such as VGG16 and ResNet50 and explored how their predictions could support pathologists during diagnosis. The findings indicated that deep learning outputs could improve the diagnostic performance of medical professionals when used as supplementary information. This research highlights the importance of combining automated image analysis with expert medical evaluation.



4. Beatriz Matias Santana Maia et al. (2024)

Maia et al. (2024) explored convolutional neural networks, transformer-based approaches and few-shot learning for classifying histopathological images of oral cancer. The study investigated different deep learning methods using an oral cancer image dataset and examined their ability to distinguish relevant tissue patterns. The research demonstrated the potential of advanced image-classification techniques for supporting oral cancer analysis. It also provides a foundation for exploring different model architectures and improving classification performance in future systems.

5. Jatender Kumar, Munish Kumar and M. K. Jindal (2026)

Kumar et al. (2026) compared four deep learning architectures—ResNet50, DenseNet201, EfficientNetB0 and ConvNeXt-Tiny—for classifying oral squamous cell carcinoma using histopathological images. The study used 10,000 images to evaluate the models for benign and malignant classification. EfficientNetB0 achieved the highest reported accuracy of 97.6% in their experimental setting. This research emphasizes the importance of selecting an appropriate deep learning architecture and evaluating models using reliable performance measures.

Comparison Table

Author & Year	Method Used	Advantages	Limitations
Das et al. (2023)	CNN and deep learning architectures such as VGG16, VGG19, ResNet50 and MobileNet	Automates oral cancer image classification and supports identification of cancer-related patterns.	Performance depends on image quality, dataset size and model training.
Panigrahi et al. (2023)	Deep transfer learning	Uses pretrained models for feature extraction and reduces the need to train models from scratch.	May be affected by dataset variability and differences between training and testing images.
Nagatsuka et al. (2023)	CNN architectures, including VGG16 and ResNet50	Can provide supplementary information to support histopathological examination.	Requires expert validation and cannot independently replace clinical diagnosis.
Maia et al. (2024)	CNNs, transformer-based methods and few-shot learning	Explores different approaches to improve oral cancer image classification.	Advanced architectures may require greater computational resources and sufficient representative data.
Kumar et al. (2026)	ResNet50, DenseNet201, EfficientNetB0 and ConvNeXt-Tiny	Compares multiple architectures and evaluates their ability to distinguish benign and malignant images.	Reported performance may not generalize to other datasets or real-world clinical settings without external validation.



V. WORKING OF SYSTEM




Fig 1: Design of the system

1. User Registration and Login

The system allows users to register by entering the required personal details. Registered users can log in using their credentials to access the application. User information is stored in an SQLite database. This process helps manage user access to the oral cancer detection system.

2. Oral Image Upload

After logging in, the user can upload an oral image through the web interface. The uploaded image is used as input for the deep learning model. The system accepts the image for further processing and classification. A simple interface makes the image submission process convenient for users.

3. Image Preprocessing

The uploaded image is resized to (224×224) pixels to match the input requirements of the MobileNetV2 model. The pixel values are normalized by dividing them by 255. This converts the pixel values into a suitable numerical range for model processing. Preprocessing helps maintain a consistent input format for image classification.

4. Deep Learning Model Processing

The system uses a Convolutional Neural Network (CNN) based on the MobileNetV2 architecture. The trained model analyzes the preprocessed image and extracts relevant visual features. These features are used to generate an output score for classification. The model performs the main image-analysis task within the application.

5. Oral Image Classification

The model produces a prediction score for the uploaded oral image. The system uses a threshold value of 0.5 to determine the predicted category. If the output score is 0.5 or higher, the image is classified as malignant; otherwise, it is classified as benign. This result represents the model's prediction and does not confirm the presence or absence of cancer.

6. Display of Prediction Results

After classification, the application displays the predicted category on the web interface. It also presents the associated confidence score to help users understand the model's output. The result provides preliminary information about the uploaded image. Users should consult qualified medical professionals for proper examination and diagnosis.

7. Data Storage and Analytics

The application uses an SQLite database to store registered user information. Supporting files, including results.json and analytics.json, are used for prediction-related and analytics information. These storage mechanisms support the



application's user management and related functions. The exact information retained depends on the implemented application logic.

8. Medical Evaluation

The system is designed to support preliminary oral cancer screening through automated image classification. It may help users recognize when further medical evaluation is necessary. However, the prediction should not be treated as a final medical diagnosis. A qualified healthcare professional must conduct the necessary clinical examination and investigations, including a biopsy when indicated.

VI. SYSTEM DESIGN

6.1 User Interface Design

The user interface provides a simple and convenient way for users to interact with the Oral Cancer Detection and Medical Diagnosis System. It allows users to register, log in, upload oral images, and view prediction results. The interface is developed using HTML, CSS, and JavaScript. It communicates with the Flask backend to process user requests. The design aims to make the application easy to use and understand.

6.2 User Authentication Module

The user authentication module manages registration and login functionality. New users can create an account by entering the required details, while registered users can log in using their credentials. User information is stored in an SQLite database. The Flask application processes authentication requests and manages access to the system. This module helps organize user access to the application.

6.3 Image Upload Module

The image upload module allows users to select and submit oral images for analysis. The uploaded image is received by the Flask application through the web interface. The image is then passed to the preprocessing stage before being supplied to the trained model. This module connects the user interface with the image classification process. The quality and suitability of the input image can affect the prediction.

6.4 Image Preprocessing Module

The preprocessing module prepares the uploaded image according to the model's input requirements. The image is resized to (224×224) pixels so that it matches the expected input dimensions of MobileNetV2. Pixel values are normalized by dividing them by 255. This process converts the image into a suitable numerical format for model processing. Consistent preprocessing helps maintain the expected input structure.

6.5 Deep Learning Model Module

The deep learning module uses a Convolutional Neural Network (CNN) based on the MobileNetV2 architecture. The trained model processes the preprocessed oral image and generates a prediction score. The application uses a threshold value of 0.5 to classify the output as benign or malignant. The model is loaded by the Flask application from the saved H5 file. The reliability of its predictions depends on the quality of the model and its training data.

6.6 Prediction Result Module

The prediction result module presents the classification generated by the deep learning model. If the output score is at least 0.5, the application classifies the image as malignant; otherwise, it classifies it as benign. The result page also displays the associated confidence score. This information helps users understand the model's output. The prediction is intended for preliminary screening and does not establish a medical diagnosis.

6.7 Database and Data Storage Module

The database module uses SQLite to store user account details, including usernames, email addresses, and password hashes. The application also uses supporting JSON files, such as results.json and analytics.json, for prediction-related and analytics information. These storage components support user management and other application functions. The exact data retained depends on the implemented application logic. Proper handling of stored information is important for maintaining user privacy.




VII. RESULTS

7.1 System Performance in Real-Time Detection

7.1 Home Page and Image Upload

The home page provides access to the oral cancer detection application and its image-upload functionality. Users can select an oral image and submit it for preliminary screening. The interface guides users through the initial stage of the classification process. This is the starting point for interacting with the system.


7.2 User Sign-In Interface



The sign-in interface features a white background with a red sidebar on the right. The main area is titled "Sign In" and includes social login buttons for Google, Facebook, and LinkedIn. Below these, it says "or use your email password". There are input fields for "Email" and "Password", a "Forgot Your Password?" link, and a red "SIGN IN" button. The sidebar has a "Hello, Friend" greeting, a "Register with your personal details to use all site features" prompt, and a "SIGN UP" button.

The sign-in interface allows registered users to enter their email address and password. It provides access to the application's image-upload and prediction features. This screen represents the user authentication stage of the system.

7.3 User Registration Interface



The registration interface features a white background with a red sidebar on the left. The main area is titled "Create Account" and includes social login buttons for Google, Facebook, and LinkedIn. Below these, it says "or use your email for". There are input fields for "Name", "Email", and "Password", and a red "SIGN UP" button. The sidebar has a "Welcome Back!" greeting, a "Enter your personal details to use all of site features" prompt, and a "SIGN IN" button.


The registration interface provides fields for creating a new user account. It supports the application's account-management functionality and works with the SQLite database used to store user information.

7.4 Oral Image Upload Interface

The upload interface allows users to select an oral image for processing. The uploaded image is passed through preprocessing before being supplied to the trained MobileNetV2 model for classification.




7.5 Analytics Dashboard



The analytics dashboard displays summary information and graphical trends related to the application's records. The screenshot demonstrates the dashboard layout; it does not independently establish the accuracy or clinical reliability of the model.

7.6 Prediction Result Interface



The prediction result interface displays the model's predicted category and associated confidence value. The classification follows the configured threshold of 0.5. The displayed output is intended for preliminary screening and must not be treated as a confirmed cancer diagnosis.

7.7 Overall Result

The supplied report documents the main screens and intended workflow of the oral cancer detection application, including account management, image uploading, prediction display, and analytics. The system integrates Flask with a trained MobileNetV2 model to generate image classifications. However, the report does not provide verified accuracy, sensitivity, specificity, or independent clinical validation results. Therefore, the results demonstrate the application's interface and workflow rather than proven diagnostic performance.

VIII. CONCLUSION

The Oral Cancer Detection and Medical Diagnosis System Using Deep Learning Techniques demonstrates the use of CNN and MobileNetV2 for preliminary oral image classification into benign and malignant categories. The system integrates Python, Flask, TensorFlow/Keras, and SQLite to provide user registration, image uploading, prediction, and result display functionalities. It offers a practical platform for exploring AI-based oral cancer screening. However, further testing and clinical validation are required to establish its reliability, and the predictions cannot replace professional medical examination or confirmed diagnosis.



IX. FUTURE SCOPE

The Oral Cancer Detection and Medical Diagnosis System Using Deep Learning Techniques can be further improved by using larger and more diverse oral image datasets to enhance the model's classification performance. Advanced deep learning architectures can be explored and compared with MobileNetV2 to identify suitable approaches for oral image analysis. Future development may include evaluating the system using performance measures such as accuracy, precision, recall, F1-score, sensitivity, and specificity. The application can also be extended into a mobile platform to improve accessibility and may support real-time image capture and analysis under suitable conditions. Additional features, such as secure patient records, improved data protection, detailed prediction reports, and integration with healthcare professionals' workflows, can also be introduced. Most importantly, the system should undergo independent testing and clinical validation with qualified medical professionals before being considered for practical healthcare use, as its predictions cannot replace professional examination or confirmed medical diagnosis.

REFERENCES

1. Panigrahi, S., Nanda, B. S., Bhuyan, R., Kumar, K., Ghosh, S., & Swarnkar, T. (2023). Classifying histopathological images of oral squamous cell carcinoma using deep transfer learning. *Heliyon*, 9(3), e13444. <https://doi.org/10.1016/j.heliyon.2023.e13444>
2. Pirayesh, Z., Mohammad-Rahimi, H., Ghasemi, N., et al. (2024). Deep learning-based image classification and segmentation on digital histopathology for oral squamous cell carcinoma: A systematic review and meta-analysis. *Journal of Oral Pathology & Medicine*, 53(9), 551–566. <https://doi.org/10.1111/jop.13578>
3. Kumar, J., Kumar, M., & Jindal, M. K. (2026). Efficient deep learning models for oral squamous cell carcinoma classification in histopathological images. *Scientific Reports*, 16, Article 21734. <https://doi.org/10.1038/s41598-026-44424-5>
4. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4510–4520). <https://doi.org/10.1109/CVPR.2018.00474>
5. Howard, A. G., Zhu, M., Chen, B., et al. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv*. <https://doi.org/10.48550/arXiv.1704.04861>
6. LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324. <https://doi.org/10.1109/5.726791>
7. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25.
8. Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations (ICLR)*. <https://doi.org/10.48550/arXiv.1409.1556>
9. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778). <https://doi.org/10.1109/CVPR.2016.90>
10. Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the International Conference on Machine Learning* (pp. 6105–6114).
11. Litjens, G., Kooi, T., Bejnordi, B. E., et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60–88. <https://doi.org/10.1016/j.media.2017.07.005>
12. Esteva, A., Kuprel, B., Novoa, R. A., et al. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542, 115–118. <https://doi.org/10.1038/nature21056>
13. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press. <https://www.deeplearningbook.org/>



14. Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 1251–1258). <https://doi.org/10.1109/CVPR.2017.195>
15. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. International Conference on Learning Representations (ICLR). <https://doi.org/10.48550/arXiv.1412.6980>
16. Abadi, M., Agarwal, A., Barham, P., et al. (2016). TensorFlow: Large-scale machine learning on heterogeneous distributed systems. arXiv. <https://doi.org/10.48550/arXiv.1603.04467>
17. TensorFlow. (n.d.). TensorFlow documentation. <https://www.tensorflow.org/>
18. Keras. (n.d.). Keras documentation. <https://keras.io/>
19. Pallets Projects. (n.d.). Flask documentation. <https://flask.palletsprojects.com/>
20. SQLite. (n.d.). SQLite documentation. <https://www.sqlite.org/docs.html>

